



Project Acronym: **OPTIMIS**
Project Title: **Optimized Infrastructure Services**
Project Number: **257115**
Instrument: **Integrated Project**
Thematic Priority: **ICT-2009.1.2 – Internet of Services, Software and Virtualisation**

Admission Control User Guide

Activity 3: Service Deployment

WP 3.5: Admission Control

Due Date:		M22
Submission Date:		31/05/2012
Start Date of Project:		01/06/2010
Duration of Project:		36 months
Organisation Responsible for the Deliverable:		ICCS/NTUA
Version:		0.1
Status		Draft, under review, final for submission
Author(s):	Kleopatra Konstanteli Sotiris Stamokostas Hassan Rasheed Juan Luis Prieto Martinez	ICCS/NTUA ICCS/NTUA SCAI ATOS
Reviewer(s)	-	-



Project co-funded by the European Commission within the Seventh Framework Programme		
Dissemination Level		
PU	Public	X
PP	Restricted to other programme participants (including the Commission)	
RE	Restricted to a group specified by the consortium (including the Commission)	
CO	Confidential, only for members of the consortium (including the Commission)	



Version History

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	2/4/2012	Compiled input	Kleopatra Konstanteli Juan Luis Prieto Martinez
0.2	3/4/2012	Compiled input	Kleopatra Konstanteli Sotiris Stamokostas Hassan Rasheed
0.3	7/4/2012	Ready for internal review	Kleopatra Konstanteli
1.0	31/5/2012	Final – ready for submission	Kleopatra Konstanteli



Table of Contents

1	INTRODUCTION	6
1.1	GLOSSARY OF ACRONYMS.....	6
2	ADMISSION CONTROL USER GUIDE	7
2.1	RELEASE INFORMATION	8
2.2	INTRODUCTION.....	8
2.3	USE CASES.....	8
2.4	SEQUENCES	9
2.4.1	<i>Single-mode operation</i>	<i>9</i>
2.4.2	<i>Federation-mode operation.....</i>	<i>11</i>
2.5	FUNCTIONALITIES.....	12
2.5.1	<i>Gatewaying and orchestration</i>	<i>12</i>
2.5.2	<i>Workload analysis</i>	<i>12</i>
2.5.3	<i>Generating input for the optimization model.....</i>	<i>13</i>
2.5.4	<i>Perform admission control test.....</i>	<i>13</i>
2.5.5	<i>Setting IP's TREC-based acceptance policy</i>	<i>13</i>
2.6	KNOWN LIMITATIONS	14
2.6.1	<i>Sensitive to interfaces changes.....</i>	<i>14</i>
2.6.2	<i>Changes in the Service Manifest schema.....</i>	<i>14</i>
2.6.3	<i>Changes in the GAMS model</i>	<i>14</i>
2.6.4	<i>GAMS license</i>	<i>14</i>
2.7	GETTING STARTED.....	14
2.7.1	<i>Contributors.....</i>	<i>15</i>
3	REFERENCES	16



Index of Figures

Figure 1: High-level view of the Admission Control architecture.	7
Figure 2: Admission Control use cases.....	9
Figure 3: Sequence diagram for single-mode operation.....	10
Figure 4: Sequence diagram for federated-mode operation	11

Index of Tables

Table 1: AC components release information.	8
Table 2: AC interfaces	14



1 Introduction

This document serves as the user guide for the software component Admission Control Y2 prototype of the OPTIMIS project, whose detailed design is included in D1.2.2.2 “Detailed Design” [5].

In more detail the document at hand focuses on:

- **Design and use cases:** A general overview of the Admission Control design, including high level objectives, and a description of the Admission Control module service-oriented architecture along with detailed sequence diagrams and use cases.
- **Functionality and interfaces:** A description of the functionality and the interfaces of each Admission Control component, the way they are incorporated into the overall OPTIMIS framework and their known limitations.

It should be noted that detailed instructions on how to install and use each AC component, including software dependencies and supported platforms as well as configuration instructions can be found in the Admission Control Installation Guide [3] and README.txt [4] and are not replicated here for the sake of brevity. Furthermore, in the Admission Control Scientific Report [2] one can find a detailed description of the optimization algorithm used behind the Admission Control Component.

1.1 Glossary of Acronyms

Acronym	Definition
AC	Admission Control
CO	Cloud Optimizer
CQoS	Cloud Quality of Service Agreement Factory
BARON	Branch and Reduce Optimization Navigator
CPU	Central Processing Unit
D	Deliverable
DRS	Document Review Sheet
EC	European Commission
GAMS	General Algebraic Modeling System
IP	Infrastructure Provider
MINLP	Mixed integer Non-Linear Programming
QoS	Quality of Service
SDO	Service Deployment Optimizer
SLA	Service Level Agreement
SP	Service Provider
TREC	Trust, Risk, Eco and Cost
VM	Virtual Machine
VMU	Virtual Machine Unit
WAR	Web ARchive
WP	Work Package
WS	Web Service



2 Admission Control User Guide

The Admission Control (AC) component is a set of four service-oriented components (RESTful services) with distinct tasks, that work together to provide the overall admission control functionality in OPTIMIS. Their tasks cover workload analysis, Trust-Risk-Eco-Cost (TREC) information retrieval, optimization modelling and admission control, and web service (WS) access to the aforementioned functionality. The services that encapsulate the capabilities mentioned above are the *Workload Analyzer*, the *TREC Analyzer*, the *Admission Controller* and the *AC Gateway* respectively.

During admission control, the Admission Controller formulates and solves a mathematical optimization problem [2][8]Error! Reference source not found. for finding the optimum allocation of the services to be deployed inside the IP's resources by communicating with the General Algebraic Modeling System (GAMS) [6]. A high-level view of the AC architecture and the way the AC components (shown in blue color) work together to provide the overall AC functionality is shown in Figure 1.

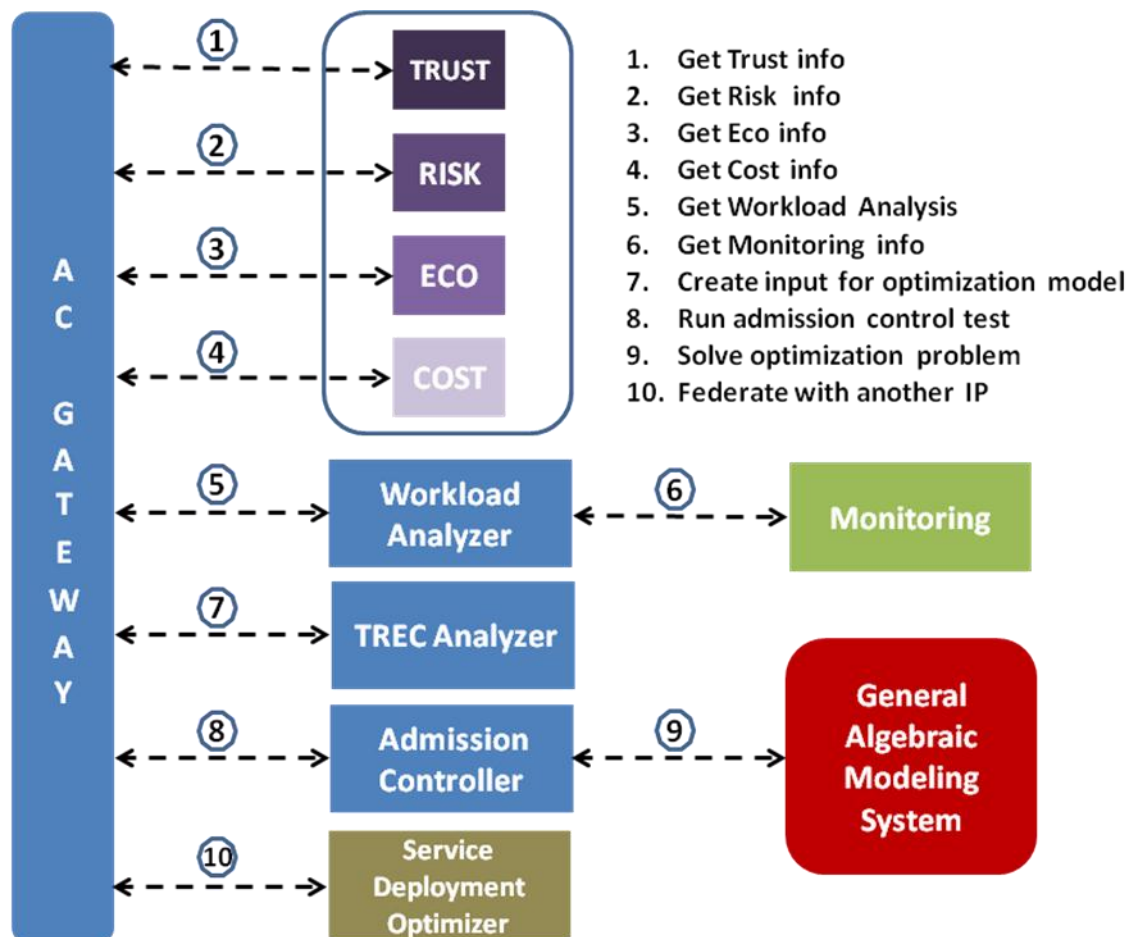


Figure 1: High-level view of the Admission Control architecture.



2.1 Release information

Component Name	Release Number	Release Date
AC Gateway	1.0-SNAPHOT	4/2012
Workload Analyzer	1.0-SNAPHOT	4/2012
TREC Analyzer	1.0-SNAPHOT	4/2012
Admission Controller	1.0-SNAPHOT	4/2012

Table 1: AC components release information.

2.2 Introduction

As already explained, the AC is a set of four RESTful components:

AC Gateway: the AC Gateway component is the Admission Control's entry point. Apart from serving as a single point of access to the overall AC functionality, the AC Gateway also performs calls to external to the AC components, as well as orchestrating the internal interactions among the rest AC components.

Workload Analyzer: the Workload Analyzer component is responsible assisting the process of admission control by providing information about the workload of the underlying Cloud infrastructure. To this direction the Workload Analyzer establishes communication with the Monitoring component for obtaining critical information regarding the current workload of the available physical resources in the Cloud, either in use or idling. The output of this analysis serves as input to the Admission Controller component.

TREC Analyzer: the TREC Analyzer is a supportive component to the Admission Controller, whose main responsibility is to transform the results of the Workload Analyzer and those of the TREC Assessment tools, as well as the requirements coming from the Service Manifest into a format understandable by the Admission Controller, which is the actual component that performs the admission control test.

Admission controller: The Admission Controller, the "brain" of the AC subsystem, is the AC subcomponent that is responsible for performing the admission control test. In case of acceptance, it provides the optimum allocation pattern (including the cost and related risk), namely the allocation offer, for each of the services in the manifest. This allocation offer is

As already described, the AC interacts with the following OPTIMIS components: Trust, Risk, Eco, Cost, Monitoring system, and Service Deployment Optimizer. Therefore, in order for the AC to be fully functional, the aforementioned components should be also deployed and up and running. Furthermore, the General Algebraic Modeling System needs to be deployed on the same physical machine that hosts the Admission Controller component.

2.3 Use Cases

When it comes to the AC component, two use cases have been identified:

1. **Single mode operation:** This use case corresponds to the Private and Multi-Cloud architecture ("All Optimis" and "Some Optimis"). In this use case, the AC component is invoked by the Cloud QoS Agreement Factory (CQoS) that is in turn invoked by the SP)through passing a complete service manifest to be considered for admission in the

underlying IP's Cloud that owns the AC component. In case the manifest cannot be admitted, the AC returns a negative answer.

2. **Federated mode operation:** this corresponds to the “Federated Cloud” and “Hybrid Cloud” architectures. Differently from the previous use case in which the result of inability to admit a new service resulted in denial of service, in this case the AC delegates the un-admitted part of the service to the Service Deployment Optimizer (SDO). The latter is responsible for finding an external IP (IP B) and returning the offer to the AC (of IP A) in order to merge them and return them to the SP level through the CQoS.

Note that the federation mode functionality is at the time of writing only supported by the AC in a stand-alone mode (AC to AC), and has not been integrated with the rest of the OPTIMIS modules.

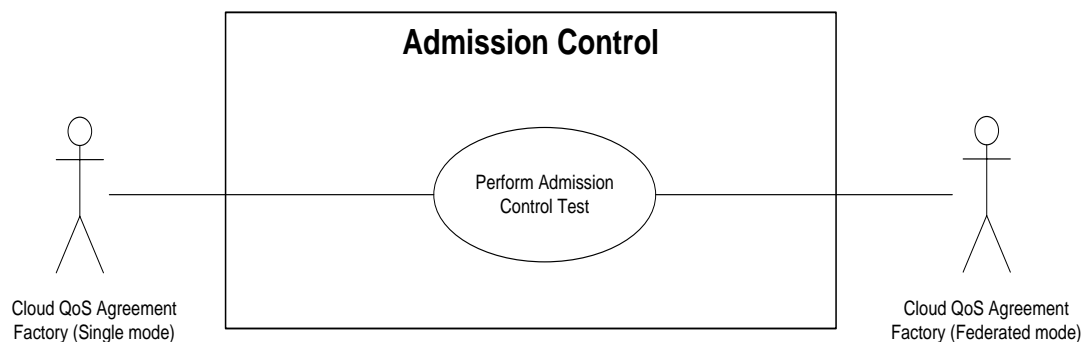


Figure 2: Admission Control use cases

2.4 Sequences

In accordance to the identified use cases, we can distinguish between two different operations of the Admission Control in OPTIMIS: (1) Single-mode operation, and (2) Federated-mode operation. Note that the 1st Year AC prototype as a whole can only handle the admission of one service (i.e. one service manifest) at a time.

2.4.1 Single-mode operation

This operation corresponds to both “Private” and “Multi-Cloud” mode, for which Admission Control will only consider the placement of all VMs inside the underlying IP Cloud. The response may as well be negative, i.e. the VMs cannot be admitted inside the given infrastructure, if Admission Control fails to accommodate at least one of them. Under this mode and in case of rejection, the AC will not consider delegation of the un-admitted part of the service manifest to other IPs. Figure 3 shows the detailed sequence diagram for this mode. In more detail the sequence involves the following steps:

- **Step 1:** CQoS triggers AC Gateway passing the service manifest of the service to be admitted, which contains the following information that is of importance to the admission control process:



- Availability constraint and reservation period
- Penalty in case of failure
- Performance requirements for all service components that comprise the service, both basic (minimum number of VMs) and elastic ones (maximum number of VMs)
- **Steps 2-5:** At next step, the AC Gateway contacts the Workload Analyzer sub-component passing it the service manifest in order to retrieve information about the current workload of the underlying infrastructure. The Workload Analyzer contacts the Monitoring component in order to retrieve monitoring data. The monitoring information is then passed onto the Workload Analyzer and the result of the workload analysis is passed back to the AC Gateway.

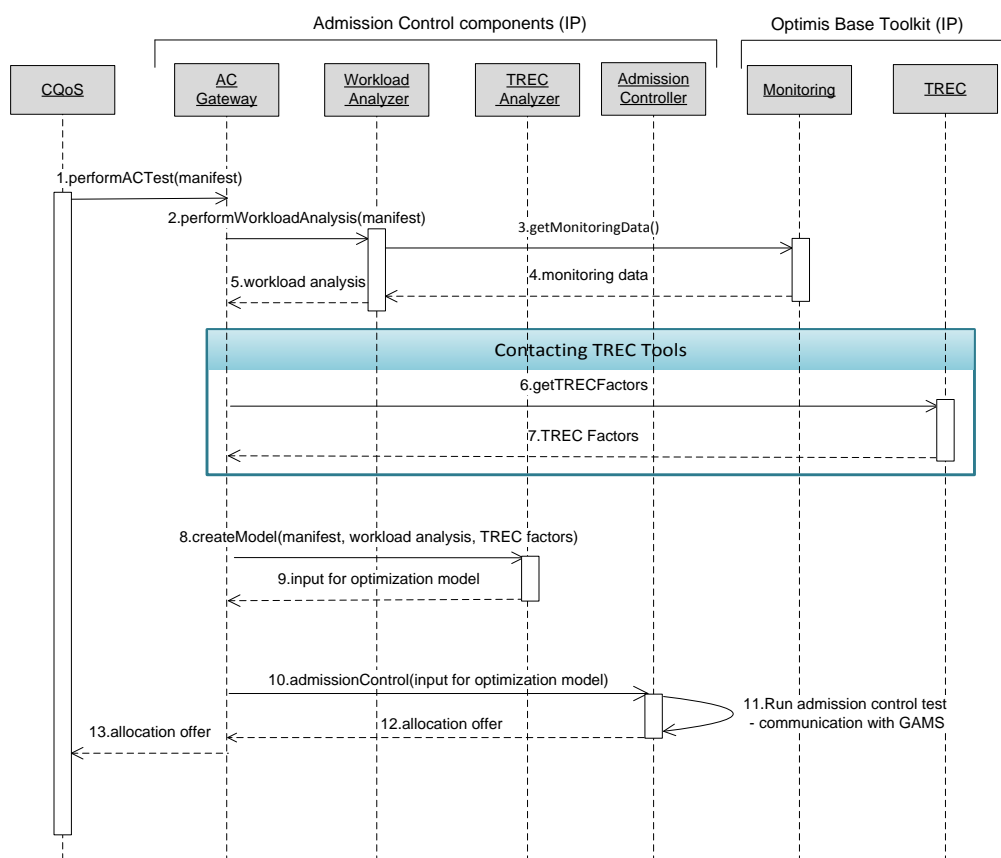


Figure 3: Sequence diagram for single-mode operation

- **Steps 6-7:** Having obtained the workload analysis, the AC Gateway contacts the TREC tools in order to retrieve information about trust, risk, eco and cost, namely the TREC factors. The AC Gateway contacts the aforementioned components through four (4) different interfaces, one for each one of them (see [3] for more details).
- **Steps 8-9:** Afterwards, the AC Gateway triggers the TREC Analyzer passing the manifest, the values of the TREC factors and the result of the workload analysis. The TREC Analyzer transforms and stores this information into several files that serve as



input to the optimization model behind the Admission Controller (see [2] for more details).

- **Steps 10-11:** The paths to the input of the optimization model along with the service manifest is then passed onto the Admission Controller, which runs the admission control test and solves the optimization problem by communicating with the General Algebraic Modeling System (GAMS).
- **Step 12:** The output of the AC in case of acceptance is an allocation offer that includes the following information:
 - Acceptance or rejection of the service, and in case of acceptance
 - Allocation offer, including the hosts that are to be used for the placement of the admitted VMs of the service.
 - The cost of the allocation, as obtained from the Cost tool.
- **Steps 13:** The allocation offer is passed to the CQoS for further assessment.

2.4.2 Federation-mode operation

Under this mode, and in the case of partial admission of the service workflow under examination, the AC may delegate the un-admitted part of the manifest to other trusted IPs, by communicating with the SDO. Upon finding an IP for the un-admitted services, the AC is responsible for merging the two allocation offers into one and passing in it back to the Cloud QoS Agreement Factory. As already stated this mode is not fully supported yet, and as the project is still evolving there is a high probability that the detailed way the AC is involved in the deployment federation is going to change. However, for completeness reasons, the way AC could operate under deployment federations is shown in Figure 4.

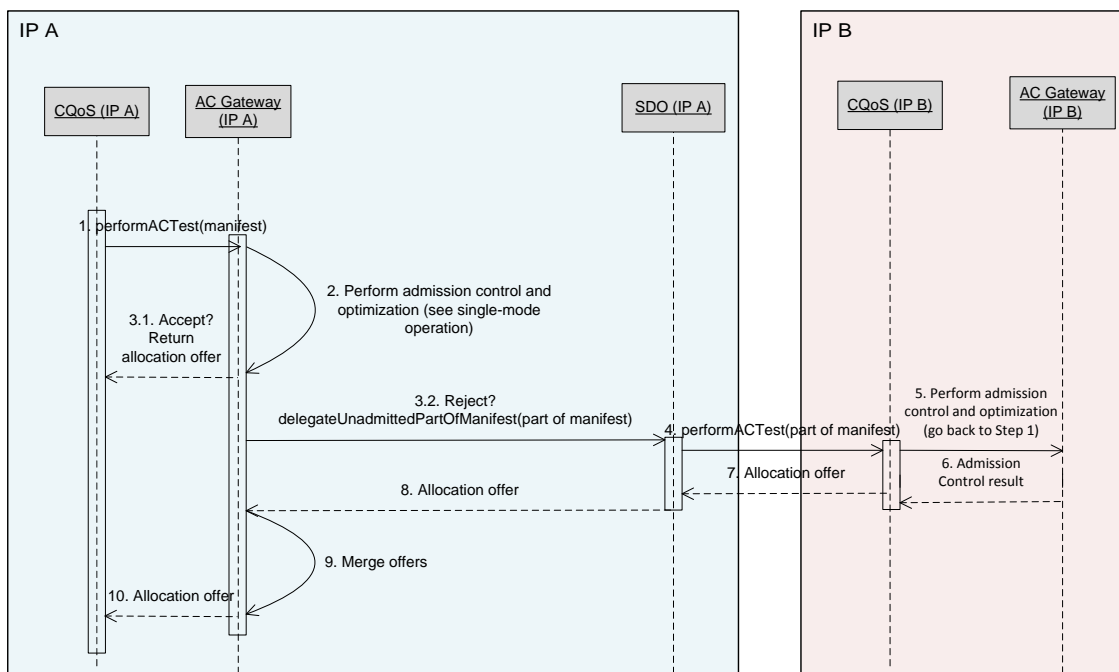


Figure 4: Sequence diagram for federated-mode operation

In more detail the sequence involves the following steps:

- **Step 1:** CQoS triggers the AC Gateway passing the service manifest to be admitted.
- **Step 2:** The AC performs all steps that are involved in the single-mode admission control operation (see sequence diagram in Figure 3).
- **Step 3.1:** In case of full acceptance, the AC Gateway returns the allocation offer to the CQoS for assessment.
- **Step 3.2:** In case the service can only be partly admitted into the infrastructure, the AC Gateway delegates the un-admitted part of the manifest to the SDO component that resides at the IP A domain.
- **Step 4-5:** The SDO processes the information and launches a new admission control test for the un-admitted part of the manifest by contacting CQoS (IP B). The later one forwards the test to the AC at another IP domain (IP B).
- **Step 6-8:** The output of the AC (IP B) is passed back to the CQoS (IP B), which passes it back to the SDO (IP A) and the later back to the AC at IP A.
- **Step 9-10:** The AC Gateway (IP A) merges the two offers and returns the result to the CQoS at IP A.

2.5 Functionalities

2.5.1 Gatewaying and orchestration

The AC Gateway implements a set of RESTful interfaces that allow other components to make use of the overall AC functionality. The responsible for triggering the aforementioned set of calls is the Cloud QoS Agreement Factory (CQoS) by passing the service manifest in order for the AC to analyze it internally. The response send back to the CQoS component is an updated version of the service manifest that includes the decision of whether the service will be admitted or not, and in case of acceptance, it also includes the optimal allocation pattern of the service components inside the Cloud physical resources (see method *performACTest* of Table 2 for more details).

It should also be noted that in case of partial acceptance of a service, and given that federation is allowed by the SP and that there are no restrictive affinity rules in place when it comes to the deployment of the service components and the service as a whole, the AC Gateway contacts the Service Deployment Optimizer (SDO) for establishing possible federation with another IP.

2.5.2 Workload analysis

The Workload Analyzer's main feature is the retrieval of monitoring data of the underlying Cloud infrastructure from the Monitoring component. Upon request from the AC Gateway, the Workload Analyzer queries the Monitoring component for the load of the available infrastructure by using the interfaces provided by the Monitoring component. It first retrieves a list of all physical hosts, and the virtual hosts running on a particular physical host. It then retrieves monitoring data sets for all physical and virtual hosts. Finally, a workload analysis document is created for all processed metrics, which is then returned to the AC Gateway and then retrieved by the TREC Analyzer component for generating the necessary by the Admission



Controller input to run the admission control test (see method *getWorkloadAnalyzer* of Table 2 for more details).

2.5.3 Generating input for the optimization model

AC Gateway triggers TREC Analyzer, once the former receives the results from both the Workload Analyzer and the TREC tools. These results are fed to the TREC Analyzer, whose main functionality is to create several input files in CVS format needed by GAMS optimization model behind the Admission Controller (see method *createModel* of Table 2 for more details).

2.5.4 Perform admission control test

In accordance to the implementation of the rest AC and OPTIMIS subcomponents, the Admission Controller's functionality is offered as a web service through a RESTful interface. The latter exposes an operation for launching the admission control test, which is triggered by the AC Gateway according to the sequence in Figure 1. In more detail, the AC Gateway contacts Admission Controller through a RESTful call, passing the path where the input files for the GAMS model are located. The Admission Controller then communicates with GAMS (which gets its input values from the CVS files created by the TREC Analyzer). The Admission Controller then parses the GAMS' output and generates a response which includes the optimum allocation scheme (in case of acceptance), which is passed back to the AC Gateway (see method *getAdmissionControl* of Table 2 for more details).

2.5.5 Setting IP's TREC-based acceptance policy

The AC through its Gateway exposes an interface for dynamically setting the weights of the TREC factors in the overall objective that is optimized. For each of the four TREC weights, a CVS file is created, which are later used inside the optimization model that runs in GAMS. This method is offered as part of the holistic management that drives decisions in OPTIMIS (see method *setPolicy* of Table 2 for more details).

NOTE: For a step-by-step explanation and examples about how to create and use a client for this method please see *Admission_Control_Installation_Guide.doc*.

Method Name	Description	Input	Output	Owner	Actor
performACTest	Perform admission control test	List of Service manifests	List of (updated) Service manifests	AC Gateway	CQoS
getWorkloadAnalyzer	Receive workload analysis	Service manifest	Workload analysis document	Workload Analyzer	AC Gateway
createModel	Generate input for optimization model	Service manifest, Workload analysis document, TREC factors	Input to optimization model (CVS files)	TREC Analyzer	AC Gateway
getAdmissionControl	Perform optimum allocation	Input to optimization model	Allocation offer	Admission Controller	AC Gateway
setPolicy	Set TREC	TREC policy	Change	AC	Cloud



	policy	levels (predetermined classes)	internal AC TREC policy	Gateway	Optimizer
--	--------	--------------------------------------	----------------------------	---------	-----------

Table 2: AC interfaces

2.6 Known limitations

These limitations are known to exist in this release of the Admission Control:

2.6.1 Sensitive to interfaces changes

Admission Control, as already mentioned, communicates with six OPTIMIS components (Trust, Risk, Eco, Cost, Monitoring system and SDO). Therefore any changes in the interfaces of the aforementioned components may result in communication errors. A good workaround is the use of updated clients for all the components involved in the communications.

2.6.2 Changes in the Service Manifest schema

Admission Control builds heavily on the service manifest. Changes in the service manifest schema therefore will imply that some of the values needed and that are extracted from the manifest will be missed or could lead to errors. To this direction, the AC should be constantly updated with the service manifest related XML beans every time the service manifest schema changes.

2.6.3 Changes in the GAMS model

Potential changes in the GAMS model may result in altering the configuration files format needed from GAMS to operate. The workaround is to make sure the configuration files have always the format expected by GAMS.

2.6.4 GAMS license

GAMS is not open-source. Without a license file GAMS can only operate in free demonstration mode with the following limitations:

1. Model limits:
 - Number of constraints and variables: 300.
 - Number of nonzero elements: 2000 (of which 1000 nonlinear).
 - Number of discrete variables: 50.
2. Additional limit for global solvers: Number of constraints and variables: 10.

The Y2 optimization model that is used to perform Admission Control exceeds the aforementioned limits, and therefore in order to solve it a GAMS base license must be acquired. This base license includes a set of MINLP (Mix Integer Non Linear Programming) solvers that could be used for solving the problem. If there is no license in place, GAMS will abort its execution and report an error message. For achieving the best performance, an extra GAMS license for the Branch and Reduce Optimization Navigator (BARON)[7] may be optionally acquired.

2.7 Getting Started

For a step-by-step guidance for installing, configuring and using each AC component please see [Admission_Control_Installation_Guide.doc](#).



Source code and information about the structure of the source directories can be found in README.txt files for each AC component.

2.7.1 Contributors

The people that contributed to the development of the AC components are listed below:

AC Gateway: Juan Luis Prieto Martinez (ATOS), Sotiris Stamokostas (ICCS/NTUA)

Workload Analyzer: Hassan Rasheed (SCAI Fraunhofer)

TREC Analyzer: Sotiris Stamokostas (ICCS/NTUA), George Katsis (ICCS/NTUA)

Admission Controller: Sotiris Stamokostas (ICCS/NTUA), Konstantinos Psychas (ICCS/NTUA), George Katsis (ICCS/NTUA), Kleopatra Konstanteli (ICCS/NTUA)



3 References

- [1] Admission Control Mechanisms Prototype D3.5.2.1 of OPTIMIS project, ICCS/NTUA and other partners, October 2011.
- [2] Admission Control Scientific Report, ICCS/NTUA and other partners, May 2012.
- [3] Admission Control Installation Guide, ICCS/NTUA and other partners, May 2012.
- [4] AdmissionControl_README.txt, ICCS/NTUA and other partners, May 2012.
- [5] D1.2.2.2 OPTIMIS Detailed Design, Umeå University and other partners, December 2011.
- [6] General Algebraic Modeling System, GAMS, <http://www.gams.com/>.
- [7] Nikolaos V. Sahinidis. BARON Branch and Reduce Optimization Navigator User's Manual v4.0. University of Illinois at Urbana-Champaign, Department of Chemical Engineering, June 2000. Available at <http://archimedes.cheme.cmu.edu/baron/manuse.pdf>.
- [8] Kleopatra Konstanteli, Tommaso Cucinotta, Konstantinos Psychas, and Theodora Varvarigou. Admission Control for Elastic Cloud Services. To appear in the Proceedings of 2012 IEEE 5th International Conference on Cloud Computing Research Track (CLOUD), <http://www.thecloudcomputing.org/2012/news.html>.